

EL764085731

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Encrypted Key Cache**

Inventor(s):

**John R. Douceur**

**Atul Adya**

**William J. Bolosky**

**Marvin M. Theimer**

ATTORNEY'S DOCKET NO. MS1-734US

**TECHNICAL FIELD**

This invention relates to cryptography, and more particularly to improving decryption performance using an encrypted key cache.

**BACKGROUND OF THE INVENTION**

File systems manage files and other data objects stored on computer systems. File systems were originally built into the computer operating system to facilitate access to files stored locally on resident storage media. As personal computers became networked, some file storage capabilities were offloaded from individual user machines to special storage servers that stored large numbers of files on behalf of the user machines. When a file was needed, the user machine simply requested the file from the server. In this server-based architecture, the file system is extended to facilitate management of and access to files stored remotely at the storage server over a network.

Today, file storage is migrating toward a model in which files are stored on various networked computers, rather than on a central storage server. One challenge faced in storing files on remote computers concerns controlling access to files that may be distributed over many different computers in a manner that allows an authorized user to access a file while at the same time insuring that unauthorized users are prevented from accessing the file. A co-pending U.S. Patent Application Serial No. \_\_\_\_\_ entitled "On-Disk File Format for Serverless Distributed File System", Attorney Docket No. MS1-733US, to inventors William J. Bolosky, Gerald Cermak, Atul Adya, and John R. Douceur describes a file format that provides such allowances and assurances. This application is hereby incorporated by reference.

1 In accordance with this file format, files are encrypted using a symmetric  
2 key, and then the symmetric key is encrypted with the public key of a  
3 public/private key pair corresponding to each user that is authorized to access the  
4 file. Thus, only the user with the correct private key is able to decrypt the  
5 symmetric key, which can then be used to decrypt the file. Although the security  
6 provided in systems employing this file format is very good, a problem that  
7 remains is that a significant amount of computational effort is involved in the use  
8 of public/private key encryption and decryption.

9 The invention addresses these problems and provides solutions to improved  
10 decryption performance that are effective for distributed file systems, and in  
11 particular, serverless distributed file systems.

## 12 **SUMMARY OF THE INVENTION**

13 An encrypted key cache is described herein.

14 In accordance with one aspect, a file that has been encrypted using a  
15 symmetric key and that has a corresponding access control entry with the  
16 symmetric key encrypted using the public key of a public/private key pair is  
17 accessed. An encrypted key cache is accessed to determine whether an access  
18 control entry to symmetric key mapping exists in the cache for the access control  
19 entry corresponding to the file. If such a mapping exists in the cache, then the  
20 mapped-to symmetric key is obtained from the cache, otherwise the encrypted  
21 symmetric key is decrypted using the private key of the public/private key pair.

22 In accordance with another aspect, an encrypted key cache includes  
23 multiple entries each of which maps an access control entry corresponding to a file  
24 in a distributed file system to a symmetric key used to encrypt that file. The  
25

1 encrypted key cache is also stored as an encrypted file within the distributed file  
2 system, allowing the entries of the cache to be used across multiple sessions. The  
3 encrypted key cache is stored in a file that can be encrypted using the private key  
4 of a public/private key pair corresponding to a user, or alternatively encrypted  
5 using a symmetric key and the symmetric key encrypted using the private key of a  
6 public/private key pair corresponding to the user and stored in an access control  
7 entry corresponding to the encrypted file.

### 8 9 **BRIEF DESCRIPTION OF THE DRAWINGS**

10 The present invention is illustrated by way of example and not limitation in  
11 the figures of the accompanying drawings. The same numbers are used  
12 throughout the figures to reference like components and/or features.

13 Fig. 1 illustrates an exemplary network environment that supports a  
14 serverless distributed file system

15 Fig. 2 illustrates logical components of an exemplary computing device that  
16 is representative of a device that participates in a distributed file system.

17 Fig. 3 illustrates an exemplary distributed file system in accordance with  
18 certain embodiments of the invention.

19 Fig. 4 illustrates an exemplary access control list corresponding to a file in  
20 accordance with certain embodiments of the invention.

21 Fig. 5 is a flowchart illustrating an exemplary process for retrieving files  
22 using an encrypted key cache in accordance with certain embodiments of the  
23 invention.

1 Fig. 6 is a flowchart illustrating an exemplary process for storing an  
2 encrypted key cache for use across multiple sessions in accordance with certain  
3 embodiments of the invention.

4 Fig. 7 illustrates a more general exemplary computer environment that can  
5 be used to implement the distributed file system.

## 6 7 **DETAILED DESCRIPTION**

8 The discussions herein assume a basic understanding of cryptography by  
9 the reader. For a basic introduction of cryptography, the reader is directed to a  
10 text written by Bruce Schneier and entitled "Applied Cryptography: Protocols,  
11 Algorithms, and Source Code in C," published by John Wiley & Sons with  
12 copyright 1994 (or second edition with copyright 1996).

13 The following discussion is directed to an encrypted key cache used to  
14 decrypt files in a distributed file system. The file system is described in the  
15 context of a symbiotic, serverless, distributed file system that runs on multiple  
16 networked computers and stores files across the computers rather than on a central  
17 server or cluster of servers. The symbiotic nature implies that the machines  
18 cooperate but do not completely trust one another. The file system does not  
19 manage the storage disk directly, but rather relies on existing file systems on local  
20 machines, such as those file systems integrated into operating systems (e.g., the  
21 Windows NT® file system).

22 While the file system is described in the context of storing "files", it should  
23 be noted that other types of storable data can be stored in the file system. The  
24 term "file" is used for discussion purposes and is intended to include data objects  
25

1 or essentially any other storage subject matter that may not be commonly  
2 characterized as a "file".

3 Fig. 1 illustrates an exemplary network environment 100 that supports a  
4 serverless distributed file system. Four client computing devices 102, 104, 106,  
5 and 108 are coupled together via a data communications network 110. Although  
6 four computing devices are illustrated, different numbers (either greater or fewer  
7 than four) may be included in network environment 100.

8 Network 110 represents any of a wide variety of data communications  
9 networks. Network 110 may include public portions (e.g., the Internet) as well as  
10 private portions (e.g., an internal corporate Local Area Network (LAN)), as well  
11 as combinations of public and private portions. Network 110 may be implemented  
12 using any one or more of a wide variety of conventional communications media  
13 including both wired and wireless media. Any of a wide variety of  
14 communications protocols can be used to communicate data via network 110,  
15 including both public and proprietary protocols. Examples of such protocols  
16 include TCP/IP, IPX/SPX, NetBEUI, etc.

17 Computing devices 102-108 represent any of a wide range of computing  
18 devices, and each device may be the same or different. By way of example,  
19 devices 102-108 may be desktop computers, laptop computers, handheld or pocket  
20 computers, personal digital assistants (PDAs), cellular phones, Internet appliances,  
21 consumer electronics devices, gaming consoles, and so forth.

22 Two or more of devices 102-108 operate to implement a serverless  
23 distributed file system (although some of them may not be operational (e.g., failed  
24 or powered-down) at any given time). The actual devices included in the  
25 serverless distributed file system can change over time, allowing new devices to

1 be added to the system and other devices to be removed from the system. Each  
2 device 102-108 that is part of the distributed file system has different portions of  
3 its mass storage device(s) (e.g., hard disk drive) allocated for use as either local  
4 storage or distributed storage. The local storage is used for data that the user  
5 desires to store on his or her local machine and not in the distributed file system  
6 structure. The distributed storage portion is used for data that the user of the  
7 device (or another device) desires to store within the distributed file system  
8 structure.

9 In the illustrated example of Fig. 1, certain devices connected to network  
10 110 have one or more mass storage devices that include both a portion used by the  
11 local machine and a portion used by the distributed file system. The amount  
12 allocated to distributed or local storage varies among the devices and can vary  
13 over time. For example, device 102 has a larger percentage allocated for a  
14 distributed system portion 120 in comparison to the local portion 122; device 104  
15 includes a distributed system portion 124 that is approximately the same size as  
16 the local portion 126; and device 106 has a smaller percentage allocated for a  
17 distributed system portion 128 in comparison to the local portion 130. The storage  
18 separation into multiple portions may occur on a per storage device basis (e.g., one  
19 hard drive is designated for use in the distributed system while another is  
20 designated solely for local use), and/or within a single storage device (e.g., part of  
21 one hard drive may be designated for use in the distributed system while another  
22 part is designated for local use). Other devices connected to network 110, such as  
23 computing device 108, may not be part of the distributed file system and thus such  
24 devices do not have any of their mass storage device(s) allocated for use by the  
25 distributed system. Hence, device 108 has only a local portion 132.

1 A distributed file system 150 operates to store one or more copies of files  
2 on different computing devices 102-106. When a new file is created by the user of  
3 a computer, he or she has the option of storing the file on the local portion of his  
4 or her computing device, or alternatively in the distributed file system. If the file  
5 is stored in the distributed file system 150, the file will be stored in the distributed  
6 system portion of the mass storage device(s) of one or more of devices 102-106.  
7 The user creating the file typically has no ability to control which device 102-106  
8 the file is stored on, nor any knowledge of which device 102-106 the file is stored  
9 on. Additionally, replicated copies of the file will typically be saved, allowing the  
10 user to subsequently retrieve the file even if one of the computing devices 102-106  
11 on which the file is saved is unavailable (e.g., is powered-down, is malfunctioning,  
12 etc.).

13 The distributed file system 150 is implemented by one or more components  
14 on each of the devices 102-106, thereby obviating the need for any centralized  
15 server to coordinate the file system. These components operate to determine  
16 where particular files are stored, how many copies of the files are created for  
17 storage on different devices, and so forth. Exactly which device will store which  
18 files depends on numerous factors, including the number of devices in the  
19 distributed file system, the storage space allocated to the file system from each of  
20 the devices, how many copies of the file are to be saved, the number of files  
21 already stored on the devices, and so on. Thus, the distributed file system allows  
22 the user to create and access files (as well as folders or directories) without any  
23 knowledge of exactly which other computing device(s) the file is being stored on.

24 The files stored by the file system are distributed among the various devices  
25 102-106 and stored in encrypted form. When a new file is created, the device on



1 which the file is being created encrypts the file prior to communicating the file to  
2 other device(s) for storage. The directory entry (e.g., the file name) for a new file  
3 is also communicated to the other device(s) for storage. Additionally, if a new  
4 folder or directory is created, the directory entry (e.g., folder name or directory  
5 name) is also communicated to the other device(s) for storage. As used herein, a  
6 directory entry refers to any entry that can be added to a file system directory,  
7 including both file names and directory (or folder) names.

8 The distributed file system 150 is designed to prevent unauthorized users  
9 from reading data stored on one of the devices 102-106. Thus, a file created by  
10 device 102 and stored on device 104 is not readable by the user of device 104  
11 (unless he or she is authorized to do so). In order to implement such security, the  
12 contents of files as well as all directory entries are encrypted, and only authorized  
13 users are given the decryption key. Thus, although device 104 may store a file  
14 created by device 102, if the user of device 104 is not an authorized user of the  
15 file, the user of device 104 cannot decrypt (and thus cannot read) either the  
16 contents of the file or its directory entry (e.g., filename).

17 Files can optionally be encrypted using a technology known as “convergent  
18 encryption”. Convergent encryption has the following two properties. First, if  
19 two or more encryptable objects are identical, then even if different encryption  
20 keys are utilized to encrypt them to provide individual cipher objects, one does not  
21 need to have access to any of the encryption keys to determine from an  
22 examination of the cipher objects that the encryptable objects are identical.  
23 Second, if two or more encryptable objects are identical but are encrypted with  
24 different encryption keys, the total space that is required to store all of the cipher  
25

1 objects is proportional to the space that is required to store a single encryptable  
2 object, plus a constant amount of storage for each distinct encryption key.

3 Generally, according to convergent encryption, a file  $F$  (or any other type of  
4 encryptable object) is initially hashed using a one-way hashing function  $h$  (e.g.,  
5 SHA, MD5, etc.) to produce a hash value  $h(F)$ . The file  $F$  is then encrypted using  
6 a symmetric cipher (e.g., RC4, RC2, etc.) with the hash value as the key, or  
7  $E_{h(F)}(F)$ . Next, read access control entries are created for each authorized user who  
8 is granted read access to the encrypted file. The access control entries are formed  
9 by encrypting the file's hash value  $h(F)$  with any number of keys  $K_1, K_2, \dots, K_m$ ,  
10 to yield  $E_{K_1}(h(F)), E_{K_2}(h(F)), \dots, E_{K_m}(h(F))$ . The keys  $K$  are randomly generated  
11 and uniquely assigned to individual users. In one implementation, each key  $K$  is  
12 the user's public key of a public/private key pair. In the illustrated example, write  
13 access control is governed by the directory server that stores the directory entry for  
14 the file and it is thus not addressed by the file format (so references to "access"  
15 within this document refer to read access unless specifically identified as another  
16 type of access). Alternatively, write access control could be implemented via  
17 access control entries in a manner analogous to the read access control discussed  
18 herein.

19 With convergent encryption, one encrypted version of the file is stored and  
20 replicated among the serverless distributed file system 150. Along with the  
21 encrypted version of the file is stored one or more access control entries depending  
22 upon the number of authorized users who have access. Thus, a file in the  
23 distributed file system 150 has the following structure:

24  
25  $[E_{h(F)}(F), \langle E_{K_1}(h(F)) \rangle, \langle E_{K_2}(h(F)) \rangle, \dots, \langle E_{K_m}(h(F)) \rangle]$

1  
2 One advantage of convergent encryption is that the encrypted file can be  
3 evaluated by the file system to determine whether it is identical to another file  
4 without resorting to any decryption (and hence, without knowledge of any  
5 encryption keys). Unwanted duplicative files can be removed by adding the  
6 authorized user(s) access control entries to the remaining file. Another advantage  
7 is that the access control entries are very small in size, on the order of bytes as  
8 compared to possibly gigabytes for the encrypted file. As a result, the amount of  
9 overhead information that is stored in each file is reduced. This enables the  
10 property that the total space used to store the file is proportional to the space that is  
11 required to store a single encrypted file, plus a constant amount of storage for each  
12 additional authorized reader of the file.

13 For more information on convergent encryption, the reader is directed to  
14 co-pending U.S. Patent Application Serial No. 09/565,821, entitled "Encryption  
15 Systems and Methods for Identifying and Coalescing Identical Objects Encrypted  
16 with Different Keys", which was filed May 5, 2000, in the names of Douceur et  
17 al., and is commonly assigned to Microsoft Corporation. This application is  
18 hereby incorporated by reference.

19 For small files, the entire file is hashed and encrypted using convergent  
20 encryption, and the resulting hash value is used as the encryption key. The  
21 encrypted file can be verified without knowledge of the key or any need to decrypt  
22 the file first. For large files, the file contents are broken into smaller blocks and  
23 then convergent encryption is applied separately to each block. For example, the  
24 file F may be segmented into "n" pages  $F^0$ - $F^{n-1}$ , where each page is a fixed size  
25 (e.g., a 4Kbyte size). Convergent encryption is then applied to the file at the block

1 level. That is, each block  $F^i$  is separately hashed using a one-way hash function  
2 (e.g., SHA, MD5, etc.) to produce a hash value  $h(F^i)$ . Each block  $F^i$  is then  
3 encrypted using a symmetric cipher (e.g., RC4, RC2, etc.) with the hash value  
4  $h(F^i)$  as the key, or  $E_{h(F^i)}(F^i)$ , resulting in an array of encrypted blocks which form  
5 the contents of the file. For more information on block-by-block encryption, the  
6 reader is directed to co-pending U.S. Patent Application Serial No. \_\_\_\_\_ entitled  
7 "On-Disk File Format for Serverless Distributed File System", Attorney Docket  
8 No. MS1-733US, to inventors William J. Bolosky, Gerald Cermak, Atul Adya, and  
9 John R. Douceur.

10 Fig. 2 illustrates logical components of an exemplary computing device 200  
11 that is representative of any one of the devices 102-106 of Fig. 1 that participate in  
12 the distributed file system 150. Computing device 200 includes a server  
13 component 202, a client component 204, a mass storage device 208, and a  
14 distributed file system interface 210. Computing device 200 also typically  
15 includes additional components (e.g., a processor), however these additional  
16 components have not been shown in Fig. 2 so as not to clutter the drawings. A  
17 more general description of a computer architecture with various hardware and  
18 software components is described below with reference to Fig. 7.

19 Mass storage device 208 can be any of a wide variety of conventional  
20 nonvolatile storage devices, such as a magnetic disk, optical disk, Flash memory,  
21 and so forth. Mass storage device 208 is partitioned into a distributed system  
22 portion and a local portion; this partitioning may change over time.

23 Computing device 200 is intended to be used in a serverless distributed file  
24 system, and as such includes both a server component 202 and client component  
25 204. Server component 202 handles requests when device 200 is responding to a

1 request involving a file or directory entry stored (or to be stored) in storage device  
2 208, while client component 204 handles the issuance of requests by device 200  
3 for files stored (or to be stored) in the distributed file system. Client component  
4 204 and server component 202 operate independent of one another. Thus,  
5 situations can arise where the serverless distributed file system 150 causes files  
6 being stored by client component 204 to be stored in mass storage device 208 by  
7 server component 202.

8 Client component 204 includes a storage and retrieval control module 220,  
9 which along with interface 210, manages access to the serverless distributed file  
10 system 150 for the creation, storage, retrieval, reading, writing, modifying, and  
11 verifying of files and directories on behalf of computing device 150. The control  
12 module 220 uses a cryptographic engine 222, a comparator 224, an encrypted key  
13 cache 226, and public/private key pair 228 when handling the encrypted files 240  
14 stored in the distributed system portion of the mass storage 208. These  
15 components are described in more detail below.

16 The server component 202 includes a distributed system control module  
17 250 and a duplication identifier 252. Distributed system control module 250  
18 manages access to the encrypted files 240. It communicates with mass storage  
19 device 208 to store and retrieve encrypted files 240. Distributed system control  
20 module 250 also maintains a record of the encrypted directory entries (not shown)  
21 in mass storage device 208 that are stored at computing device 200 (or  
22 alternatively that are stored elsewhere in the serverless distributed file system).

23 Duplication identifier 252 helps identify identical encrypted files in the  
24 distributed file system. When the duplication identifier 252 finds a duplication  
25 that is not an intentional replication for fault tolerant purposes, the duplication

1 identifier 252 notifies the control module 250, which then eliminates the  
2 duplicated file and adds the access control entries to the eliminated file to the  
3 remaining file.

4 Fig. 3 illustrates an exemplary distributed file system in accordance with  
5 certain embodiments of the invention. Two computing devices 300 and 302 are  
6 illustrated, each of which is a computing device 200 of Fig. 2 that is part of  
7 distributed file system 150 of Fig. 1. For ease of explanation and to avoid  
8 cluttering the drawings, some components and modules of computing devices 300  
9 and 302 are not illustrated in Fig. 3.

10 Computing device 300 includes a storage device 208 and an interface 210.  
11 Storage device 208 can store multiple files, one of which (encrypted file 304) is  
12 illustrated in Fig. 3. File 304 has a corresponding access control list 306 including  
13 multiple access control entries 308. Files are stored on storage device 208 using a  
14 file format that is composed of two parts: a data stream and a metadata stream.  
15 The data stream contains the contents of the file, which makes up the bulk of the  
16 entire file. The data stream is a primary (unnamed) data stream, while the  
17 metadata stream of the file is a separate named stream. In the illustrated example,  
18 encrypted file 304 is the data stream and access control list is included in the  
19 metadata stream.

20 Although access control list 306 is illustrated as being stored in storage  
21 device 208 along with its corresponding file 304, access control list 306 may  
22 alternatively be stored separately (either on computing device 300 or alternatively  
23 another computing device (not shown)). For example, each file stored in storage  
24 device 208 has a corresponding directory entry in a directory service on device  
25

300 or some other device, and access control list 306 may be stored along with this corresponding directory entry.

An access control entry 308 includes an identification of the user (e.g., by the user name assigned to the user for the serverless distributed file system) and also an encrypted version of the symmetric key used to encrypt the file. The symmetric key is encrypted using the public key of a public/private key pair of the user that corresponds to the access control entry, so that only that user (who holds the private key of the public/private key pair) is able to decrypt the encrypted symmetric key from that access control entry. Thus, each access control entry will include a different encrypted symmetric key, but all will decrypt (using the appropriate private keys) to the same symmetric key.

Computing device 302 is illustrated including interface 210, storage/retrieval control module 220, cryptographic engine 222, comparator 224, encrypted key cache 226, and public/private key pair 228. Computing devices 300 and 302 communicate with one another via their respective interfaces 210. When a file is requested (e.g., by a user of computing device 302), control module 220 obtains the requested file and access control entry corresponding to the user requesting the file and communicates with cryptographic engine 222, and possibly comparator 224, to decrypt the received file.

Cryptographic engine 222 can use any one or more of a wide variety of conventional encryption algorithms (e.g., RC2, RC4, etc.) to decrypt data. Cryptographic engine 222 makes use of a symmetric key, received from another device, to decrypt data (e.g., in the discussion above regarding convergent encryption, the hash value  $h(F)$  is the symmetric key). Cryptographic engine 222 may also make use of one or both keys of public/private key pair 228, such as to

1 decrypt data that was encrypted by another device using the public key of key pair  
2 228. In certain situations, rather than having decryption module 316 use the  
3 private key of key pair 228 to perform certain decryptions, control module 220  
4 may obtain the appropriate decrypted symmetric key from a cache 226 of  
5 encrypted keys (referred to as the encrypted key cache), as discussed in more  
6 detail below.

7 To grant access privileges to multiple users, the file system 150 maintains a  
8 user key list (referred to as an access control list or ACL) for each file. Each entry  
9 (referred to as an access control entry or ACE) in the ACL contains the data used  
10 by specific users to decrypt the file. More particularly, when a file is created by a  
11 computing device, it is encrypted using a symmetric key. The symmetric key can  
12 be generated in any of a wide variety of manners, such as randomly or derived by  
13 the file contents in some manner (e.g., using convergent encryption discussed  
14 above). Separate ACLs may be maintained for each type of access right (e.g., read  
15 access, modify/write access, etc.), or alternatively a single ACL may be  
16 maintained with each ACE indicating what access rights the corresponding user  
17 has to the file. Generally, read ACLs are addressed in this document.

18 Fig. 4 illustrates an exemplary access control list corresponding to a file in  
19 accordance with certain embodiments of the invention. Each entry in list 306  
20 includes a user name 402 of the user (or user group) with access privileges and an  
21 encrypted symmetric key 404 encrypted with the public key corresponding to that  
22 user (or user group). If the file is publicly readable, the key list contains only a  
23 special entry 406, in which the FILE\_USER\_NAME is the reserved value  
24 USER\_EVERYONE to indicate that everyone has access to the file, and the  
25



1 associated encrypted key segment 404 contains the symmetric key in decrypted  
2 form.

3 Since files are stored on machines that are not trusted, read access cannot  
4 be sufficiently controlled merely by listing the authorized readers of a file in the  
5 metadata, as is commonly done in trusted file systems. Thus, this file format relies  
6 on cryptography to provide access security. Only a user with knowledge of an  
7 appropriate private key will be able to recover the symmetric key used to encrypt  
8 the file. As a result, an impostor who attempts to recover the symmetric key using  
9 an authorized user's name will be unable to decrypt the symmetric key because  
10 that impostor does not have knowledge of the user's private key.

11 Returning to Fig. 3, assume for purposes of explanation that computing  
12 device 302 desires to access encrypted file 304 (e.g., in response to a request from  
13 a user to open or print file 304). The file request is communicated to computing  
14 device 300, which in turn communicates encrypted file 304 and corresponding  
15 ACL 306 to computing device 302. Computing device 300 may return the entire  
16 ACL 306 to computing device 302, or alternatively only the ACEs that the  
17 requesting user is part of (e.g., ACEs identifying the user name or a user group  
18 that the user is part of).

19 Assuming that the entire ACL 306 is returned to computing device 302,  
20 control module 220 identifies each ACE in ACL 306 that the user of computing  
21 device 302 corresponds to (e.g., is either that user or the user is part of the group  
22 identified in the ACE). Control module 220 then has comparator 224 check  
23 encrypted key cache 226 to determine whether an ACE to symmetric key mapping  
24 exists in cache 226 for any of the ACEs that the user corresponds to. If such a  
25 mapping exists, then control module 220 forwards the mapped-to symmetric key

1 from cache 226 to cryptographic engine 222 to use to decrypt file 304. Thus, if  
2 the ACE to symmetric key mapping exists in cache 226, cryptographic engine 222  
3 does not need to expend the effort to decrypt the encrypted symmetric key from  
4 the ACE using the private key from key pair 228.

5 However, if the ACE to symmetric key mapping does not exist in cache  
6 226, then control module 220 has cryptographic engine 222 decrypt the encrypted  
7 symmetric key from the ACE using the private key from key pair 228. Although  
8 only a single key pair 228 is illustrated, computing device 302 may include  
9 multiple key pairs (one corresponding to the user's identifier and one for each  
10 additional group that the user is part of). Once decrypted, cryptographic engine  
11 222 communicates the decrypted symmetric key to control module 220, which in  
12 turn stores both the ACE and the decrypted symmetric key (and an indication that  
13 they map to one another) into cache 226.

14 Encrypted key cache 226 maintains multiple ACE to symmetric key  
15 mappings 310 (each mapping is also referred to as an entry of cache 226). Each  
16 mapping 310 includes an ACE portion 312 and a symmetric key portion 314.  
17 Each mapping 310 corresponds to an ACE corresponding to a file that has been (or  
18 can be) accessed by a user of computing device 302. Each ACE portion 312  
19 includes an ACE that corresponds to both the user of computing device 302 and  
20 the file to which mapping 310 corresponds (e.g., one of ACEs 308 if mapping 310  
21 corresponds to encrypted file 304). Each symmetric key portion 314 includes the  
22 decrypted symmetric key used to decrypt the file to which mapping 310  
23 corresponds.

24 The mapping of an ACE to decrypted symmetric key in encrypted key  
25 cache 226 can be explicit (e.g., a pointer to where the symmetric key portion of an

entry is maintained) or alternatively inherent in the data structure used (e.g., a list of (ACE, key) pairs). As discussed above, each ACE includes a user name and an encrypted symmetric key. In one implementation, encrypted key cache 226 is indexed by the encrypted symmetric key, although alternatively encrypted key cache 226 can be indexed by the user name.

Encrypted key cache 226 can be a variable size. For example, each entry 310 may be maintained in cache 226 so long as the file to which the entry 310 corresponds is requested by a component or module of device 302 within a certain time frame (e.g., at least once every seven days, at least once out of every 300 file accesses, etc.). If the file to which the entry 310 corresponds is not accessed within that certain time frame, then the entry is removed from cache 226. Thus, the size of cache 226 can change over time.

Alternatively, encrypted key cache 226 may be a fixed size. For example, encrypted key cache 226 may be limited to storing 512 entries 310. If an attempt is made to store a new entry to cache 226 and cache 226 is already full (e.g., already stores 512 entries), then a cache replacement process is used to determine which current cache entry is to be deleted to make room for the new cache entry. Any of a wide variety of conventional cache replacement processes can be used, such as the well known least-recently-used (LRU) or least-frequently used (LFU) cache replacement algorithms.

Alternatively, rather than having a timing out or LRU-based cache entry removal, an "active" invalidation model could be used. For example, a cache entry invalidation command or message could be sent by some other entity (e.g., a computing device 200 of Fig. 2) at an appropriate time, such as after the entity learns about the change of the symmetric key for that entry.

1        Encrypted key cache 226 is indexed based on the ACE portions 312. In  
2 order to determine whether a mapping for a received ACE is included in key cache  
3 226, control module 220 has comparator 224 compare the received ACE to the  
4 portions 312 and check whether any of the portions 312 match (are the same as)  
5 the received ACE. If the ACE portion 312 of one of the mappings 310 matches  
6 the received ACE, then that mapping is used to determine the symmetric key for  
7 the file.

8        Checking whether the received ACE matches any of the portions 312 can  
9 be performed in any of a wide variety of manners. For example, each portion 312  
10 may be compared to the received ACE until either a match is found or all portions  
11 312 are compared. Alternatively, any of a wide variety of conventional rapid-  
12 retrieval data structures may be used to more quickly determine whether a  
13 received ACE matches one of the portions 312. For example, a hashing table may  
14 be created (by control module 220 or alternatively another module) and the  
15 received ACE used in a conventional manner to index into the hashing table, with  
16 the resultant hash value indicating which of the portions 312 are possible matches  
17 to the received ACE and only those indicated portions 312 compared to the  
18 received ACE.

19        In one implementation, encrypted key cache 226 is also stored by  
20 computing device 302 as an encrypted file within the distributed file system 150.  
21 By so storing cache 226, cache 226 can persist across multiple sessions (each  
22 "session" refers to a user being logged-on to a computing device – a new session  
23 begins each time the user logs in to a computing device and the session ends when  
24 the user logs off the computing device (either explicitly by selecting a "log off"  
25 option or implicitly by powering-down the computing device, having the

1 computing device fail (e.g., crash), etc.)), without placing unencrypted copies of  
2 the symmetric keys contained in the cache onto persistent storage, where they  
3 might be read by unauthorized parties, e.g. after the user has logged out and  
4 another user has logged in.

5 In order to store key cache 226, when the session ends a component of  
6 computing device 302 (e.g., control module 220) stores cache 226 as a file and  
7 then has cryptographic engine 222 encrypt the file using a symmetric key. An  
8 ACE is then created for the encrypted file, the symmetric key is encrypted using  
9 the public key of the user whose session it is, and then the encrypted symmetric  
10 key is stored in the ACE. The encrypted file and corresponding ACE is then  
11 stored in the file system 150 as if it were any other file. Subsequently, when the  
12 user logs on to device 302 again (or alternatively another device), the encrypted  
13 file is retrieved, decrypted using the user's private key, and then the populated  
14 encrypted key cache 226 is available for decrypting symmetric keys of other files,  
15 including all of the entries 310 that were included in cache 226 when the user's  
16 previous session ended. It should be noted that the encrypted file containing cache  
17 226 can be retrieved as part of the log-on process, or alternatively only if needed  
18 (e.g., the first time the user attempts a read access to another file).

19 It should also be noted that the file containing cache 226 remains securely  
20 stored (in encrypted form) at the remote system even if the device 302 using cache  
21 226 to decrypt symmetric keys is improperly shut down or fails. However,  
22 changes made to the cache 226 after device 302 loaded cache 226 but before  
23 device 302 shut down or failed may not be reflected in the securely stored cache  
24 226. Alternatively, device 302 may periodically store an encrypted file containing  
25

1 cache 226 at the remote location to guard against such data loss in the event of a  
2 failure or improper shut-down.

3 Alternatively, rather than encrypting the file containing cache 226 with a  
4 symmetric key and then encrypting the symmetric key with the public key of the  
5 user, the file containing cache 226 may be encrypted just with the public key of  
6 the user (no symmetric key being used). Typically, the user (and/or distributed  
7 file system designer) will desire the file containing cache 226 to be accessible only  
8 to that user and not to any other users. Thus, by encrypting the file with the user's  
9 public key, the user will be able to decrypt the file and obtain access to populated  
10 cache 226, but others will not be able to decrypt the file.

11 In another alternative, encrypted key cache 226 may be cleared when  
12 computing device 302 is powered-down (either intentionally by deleting each  
13 entry or inherently by storing cache 226 in volatile memory). When computing  
14 device 302 is subsequently powered-on, encrypted key cache 226 is then re-  
15 populated as files are accessed by device 302.

16 Additionally, situations can arise where the symmetric key with which a  
17 file is encrypted is changed. If a change is made to the symmetric key with which  
18 a file is encrypted, all ACEs for that file will also change, since each ACE  
19 contains the encryption of the symmetric key with a user's public key. If  
20 computing device 302 accesses file 304 after the symmetric key for that file has  
21 been changed, then the new ACE 308 for the file will not be resident in the key  
22 cache 226, even if the file had recently been accessed by computing device 302.  
23 Therefore, control module 220 has cryptographic engine 222 decrypt the new  
24 encrypted symmetric key from the new ACE using the private key from key pair  
25 228. Once decrypted, the new ACE and corresponding symmetric key are stored

1 in the key cache 226 for future reference. Since the old symmetric key has been  
2 discarded, the corresponding old ACE will not be looked up in the key cache  
3 again, so the entry in the key cache that contains this ACE will eventually be  
4 removed, either when the time frame expires (for a variable size cache) or when  
5 the cache replacement algorithm removes this entry after the cache size limit has  
6 been reached (for a fixed size cache).

7 Alternatively, in some implementations it may not be readily apparent if an  
8 ACE has changed since it was inserted in key cache 226 (for example, if key cache  
9 226 is indexed by user name rather than by encrypted symmetric key). These  
10 situations can be resolved by initially attempting to decrypt the file using the  
11 symmetric key from cache 226. If the decryption is verified (e.g., using  
12 convergent encryption discussed above), then the file has been successfully  
13 decrypted. However, if the decryption fails, the correct symmetric key is  
14 decrypted from the received ACE.

15 It should be noted that, although two different computing devices 300 and  
16 302 are illustrated in Fig. 3, due to the nature of the serverless distributed file  
17 system, an encrypted file being retrieved by a computing device may actually be  
18 stored in a distributed storage portion of that device (perhaps unbeknownst to the  
19 component or module that is requesting the file). In these situations, the file is still  
20 stored in encrypted form and is decrypted upon retrieval in the same manner as  
21 discussed above.

22 Fig. 5 is a flowchart illustrating an exemplary process for retrieving files  
23 using an encrypted key cache in accordance with certain embodiments of the  
24 invention. The process of Fig. 5 is carried out by a client component 204 of  
25 Fig. 2, and may be implemented in software.

Initially, a file access request is received (act 502). In response to the request, an access control entry corresponding to the requested file and the user of the computing device making the request is obtained (act 504). A check is then made as to whether the encrypted key cache is currently loaded into the memory of the computing device (act 506). If the encrypted key cache is not loaded into memory, then the encrypted key cache is retrieved (in file form) from the location it is stored at within the distributed file system (act 508). The file containing the encrypted key cache is then decrypted based on the user's private key (act 510) and then a check made as to whether the access control entry obtained in act 504 is in the encrypted key cache (act 512). Returning to act 506, if the encrypted key cache is already loaded into memory, then processing proceeds to act 512 (bypassing acts 508 and 510).

If the access control entry obtained in act 504 is in the encrypted key cache, then the symmetric key associated with the access control entry is retrieved from the encrypted key cache (act 514) and used to decrypt the requested file (act 516). However, if the access control entry obtained in act 504 is not in the encrypted key cache, then the encrypted symmetric key included in the access control entry is decrypted using the requesting user's private key (act 518). A mapping of the access control entry to decrypted symmetric key is then stored in the encrypted key cache (act 520) and the decrypted symmetric key is used to decrypt the requested file (act 516).

Fig. 6 is a flowchart illustrating an exemplary process for storing an encrypted key cache for use across multiple sessions in accordance with certain embodiments of the invention. The process of Fig. 6 is carried out by a client component 204 of Fig. 2, and may be implemented in software.



1 Initially, the encrypted key cache is stored as a file (act 602). The file  
2 containing the encrypted key cache is then encrypted using a symmetric key (act  
3 604) and the symmetric key is encrypted using a private key of a public/private  
4 key pair corresponding to the current user of the computing device (act 606). An  
5 access control entry corresponding to the file is then created (act 608), and the file  
6 and corresponding access control entry are then stored in the distributed file  
7 system (act 610).

8 Fig. 7 illustrates a more general exemplary computer environment 700,  
9 which can be used to implement the distributed file system. The computer  
10 environment 700 is only one example of a computing environment and is not  
11 intended to suggest any limitation as to the scope of use or functionality of the  
12 computer and network architectures. Neither should the computer environment  
13 700 be interpreted as having any dependency or requirement relating to any one or  
14 combination of components illustrated in the exemplary computer environment  
15 700.

16 Computer environment 700 includes a general-purpose computing device in  
17 the form of a computer 702. Computer 702 can be, for example, any of computing  
18 devices 102-108 of Fig. 1, a computing device 200 of Fig. 2, or any of computing  
19 devices 300 or 302 of Fig. 3. The components of computer 702 can include, but  
20 are not limited to, one or more processors or processing units 704, a system  
21 memory 706, and a system bus 708 that couples various system components  
22 including the processor 704 to the system memory 706.

23 The system bus 708 represents one or more of any of several types of bus  
24 structures, including a memory bus or memory controller, a peripheral bus, an  
25 accelerated graphics port, and a processor or local bus using any of a variety of

1 bus architectures. By way of example, such architectures can include an Industry  
2 Standard Architecture (ISA) bus, a Micro Channel Architecture (MCA) bus, an  
3 Enhanced ISA (EISA) bus, a Video Electronics Standards Association (VESA)  
4 local bus, and a Peripheral Component Interconnects (PCI) bus also known as a  
5 Mezzanine bus.

6 Computer 702 typically includes a variety of computer readable media.  
7 Such media can be any available media that is accessible by computer 702 and  
8 includes both volatile and non-volatile media, removable and non-removable  
9 media.

10 The system memory 706 includes computer readable media in the form of  
11 volatile memory, such as random access memory (RAM) 710, and/or non-volatile  
12 memory, such as read only memory (ROM) 712. A basic input/output system  
13 (BIOS) 714, containing the basic routines that help to transfer information  
14 between elements within computer 702, such as during start-up, is stored in ROM  
15 712. RAM 710 typically contains data and/or program modules that are  
16 immediately accessible to and/or presently operated on by the processing unit 704.

17 Computer 702 may also include other removable/non-removable,  
18 volatile/non-volatile computer storage media. By way of example, Fig. 7  
19 illustrates a hard disk drive 716 for reading from and writing to a non-removable,  
20 non-volatile magnetic media (not shown), a magnetic disk drive 718 for reading  
21 from and writing to a removable, non-volatile magnetic disk 720 (e.g., a "floppy  
22 disk"), and an optical disk drive 722 for reading from and/or writing to a  
23 removable, non-volatile optical disk 724 such as a CD-ROM, DVD-ROM, or other  
24 optical media. The hard disk drive 716, magnetic disk drive 718, and optical disk  
25 drive 722 are each connected to the system bus 708 by one or more data media

1 interfaces 726. Alternatively, the hard disk drive 716, magnetic disk drive 718,  
2 and optical disk drive 722 can be connected to the system bus 708 by one or more  
3 interfaces (not shown).

4 The disk drives and their associated computer-readable media provide non-  
5 volatile storage of computer readable instructions, data structures, program  
6 modules, and other data for computer 702. Although the example illustrates a  
7 hard disk 716, a removable magnetic disk 720, and a removable optical disk 724,  
8 it is to be appreciated that other types of computer readable media which can store  
9 data that is accessible by a computer, such as magnetic cassettes or other magnetic  
10 storage devices, flash memory cards, CD-ROM, digital versatile disks (DVD) or  
11 other optical storage, random access memories (RAM), read only memories  
12 (ROM), electrically erasable programmable read-only memory (EEPROM), and  
13 the like, can also be utilized to implement the exemplary computing system and  
14 environment.

15 Any number of program modules can be stored on the hard disk 716,  
16 magnetic disk 720, optical disk 724, ROM 712, and/or RAM 710, including by  
17 way of example, an operating system 726, one or more application programs 728,  
18 other program modules 730, and program data 732. Each of such operating  
19 system 726, one or more application programs 728, other program modules 730,  
20 and program data 732 (or some combination thereof) may implement all or part of  
21 the resident components that support the distributed file system.

22 A user can enter commands and information into computer 702 via input  
23 devices such as a keyboard 734 and a pointing device 736 (e.g., a "mouse").  
24 Other input devices 738 (not shown specifically) may include a microphone,  
25 joystick, game pad, satellite dish, serial port, scanner, and/or the like. These and

1 other input devices are connected to the processing unit 704 via input/output  
2 interfaces 740 that are coupled to the system bus 708, but may be connected by  
3 other interface and bus structures, such as a parallel port, game port, or a universal  
4 serial bus (USB).

5 A monitor 742 or other type of display device can also be connected to the  
6 system bus 708 via an interface, such as a video adapter 744. In addition to the  
7 monitor 742, other output peripheral devices can include components such as  
8 speakers (not shown) and a printer 746 which can be connected to computer 702  
9 via the input/output interfaces 740.

10 Computer 702 can operate in a networked environment using logical  
11 connections to one or more remote computers, such as a remote computing device  
12 748. By way of example, the remote computing device 748 can be a personal  
13 computer, portable computer, a server, a router, a network computer, a peer device  
14 or other common network node, and the like. The remote computing device 748 is  
15 illustrated as a portable computer that can include many or all of the elements and  
16 features described herein relative to computer 702.

17 Logical connections between computer 702 and the remote computer 748  
18 are depicted as a local area network (LAN) 750 and a general wide area network  
19 (WAN) 752. Such networking environments are commonplace in offices,  
20 enterprise-wide computer networks, intranets, and the Internet.

21 When implemented in a LAN networking environment, the computer 702 is  
22 connected to a local network 750 via a network interface or adapter 754. When  
23 implemented in a WAN networking environment, the computer 702 typically  
24 includes a modem 756 or other means for establishing communications over the  
25 wide network 752. The modem 756, which can be internal or external to computer

1 702, can be connected to the system bus 708 via the input/output interfaces 740 or  
2 other appropriate mechanisms. It is to be appreciated that the illustrated network  
3 connections are exemplary and that other means of establishing communication  
4 link(s) between the computers 702 and 748 can be employed.

5 In a networked environment, such as that illustrated with computing  
6 environment 700, program modules depicted relative to the computer 702, or  
7 portions thereof, may be stored in a remote memory storage device. By way of  
8 example, remote application programs 758 reside on a memory device of remote  
9 computer 748. For purposes of illustration, application programs and other  
10 executable program components such as the operating system are illustrated herein  
11 as discrete blocks, although it is recognized that such programs and components  
12 reside at various times in different storage components of the computing device  
13 702, and are executed by the data processor(s) of the computer.

14 Computer 702 typically includes at least some form of computer readable  
15 media. Computer readable media can be any available media that can be accessed  
16 by computer 702. By way of example, and not limitation, computer readable  
17 media may comprise computer storage media and communication media.  
18 Computer storage media includes volatile and nonvolatile, removable and non-  
19 removable media implemented in any method or technology for storage of  
20 information such as computer readable instructions, data structures, program  
21 modules or other data. Computer storage media includes, but is not limited to,  
22 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,  
23 digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic  
24 tape, magnetic disk storage or other magnetic storage devices, or any other media  
25 which can be used to store the desired information and which can be accessed by

1 computer 702. Communication media typically embodies computer readable  
2 instructions, data structures, program modules or other data in a modulated data  
3 signal such as a carrier wave or other transport mechanism and includes any  
4 information delivery media. The term "modulated data signal" means a signal that  
5 has one or more of its characteristics set or changed in such a manner as to encode  
6 information in the signal. By way of example, and not limitation, communication  
7 media includes wired media such as wired network or direct-wired connection,  
8 and wireless media such as acoustic, RF, infrared and other wireless media.  
9 Combinations of any of the above should also be included within the scope of  
10 computer readable media.

11 The invention has been described herein in part in the general context of  
12 computer-executable instructions, such as program modules, executed by one or  
13 more computers or other devices. Generally, program modules include routines,  
14 programs, objects, components, data structures, etc. that perform particular tasks  
15 or implement particular abstract data types. Typically the functionality of the  
16 program modules may be combined or distributed as desired in various  
17 embodiments.

18 For purposes of illustration, programs and other executable program  
19 components such as the operating system are illustrated herein as discrete blocks,  
20 although it is recognized that such programs and components reside at various  
21 times in different storage components of the computer, and are executed by the  
22 data processor(s) of the computer.

23 Alternatively, the invention may be implemented in hardware or a  
24 combination of hardware, software, and/or firmware. For example, one or more  
25

1 application specific integrated circuits (ASICs) could be designed or programmed  
2 to carry out the invention.

3 It should be noted that, although discussed primarily herein with reference  
4 to a serverless distributed file system, the invention can be used in any file system  
5 in which it is desired to keep files encrypted so that various other parties cannot  
6 see the content of the files. Thus, the invention can be used in other embodiments,  
7 such as, for example, those with one or more centralized file servers.

### 8 9 **Conclusion**

10 Although the description above uses language that is specific to structural  
11 features and/or methodological acts, it is to be understood that the invention  
12 defined in the appended claims is not limited to the specific features or acts  
13 described. Rather, the specific features and acts are disclosed as exemplary forms  
14 of implementing the invention.